Journal of Monetary Economics ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at SciVerse ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Comparative advantage and unemployment

Mark Bils [a,b], Yongsung Chang [a,c,*], Sun-Bin Kim [c]

[a] University of Rochester, USA
[b] NBER, USA
[c] Yonsei University, South Korea

ARTICLE INFO

Article history:
Received 28 October 2010
Received in revised form
5 January 2012
Accepted 11 January 2012

ABSTRACT

Worker heterogeneity in productivity and labor supply is introduced into a matching model. Workers who earn high wages and work high-hours are identified as those with strong market comparative advantage—high rents from being employed. The model is calibrated to match separation, job finding, and employment in the SIPP data. The model predicts a big drop in employment for workers with weak comparative advantage during recessions. But the data show that workers with strong comparative advantage also display sizable employment fluctuations, implying that aggregate employment fluctuations are not explained by the responses of workers with small rents to employment.

1. Introduction

Can the Diamond–Mortensen–Pissarides (DMP) model capture fluctuations in unemployment? The answer depends largely on the size of rents to employment. If workers collect little surplus from being employed, then aggregate shocks can translate into large model fluctuations in vacancies and unemployment. (See Mortensen and Nagypal, 2007; Costain and Reiter, 2008; Hagedorn and Manovskii, 2008, for illustrations.) Establishing the size of these rents is difficult because they reflect hard-to-measure individual valuations of leisure and home production. But given that workers differ markedly in both their hours and earnings when working, they should differ in rents from employment—with lower rents for workers with low-hours and earnings. In turn, workers with low-hours and earnings should display larger business cycle fluctuations. These predictions are tested for workers in the Survey of Income and Program Participation (SIPP).

To do so, we exploit a Mortensen and Pissarides (1994) model with endogenous separation. The model is extended to allow for differences in workers' rents to employment (market comparative advantage) driven by heterogeneity in workers' market human capital and in their valuation of non-market time. Dispersion in workers' human capital is chosen to match the cross-sectional distribution of wage rates in the SIPP. Similarly, dispersion in workers' values of non-market time is introduced to capture the distribution of hours worked, conditional on being employed, in the SIPP. To achieve this latter mapping to the data, an intensive margin for labor supply is introduced. Workers with a high market human capital relative to non-market predictably work more hours. This identifies these workers as those with strong comparative advantage in the workforce—that is, high rents from employment.

The model of hours worked, separations, and vacancy creation is presented in Section 2. Section 3 describes the SIPP data and presents average patterns in employment, hours and turnover rates for four distinct groups of workers, stratifying both on workers' long-term wage rates and hours worked. Section 4 calibrates the model to the same four distinct groups

* Corresponding author at: Department of Economics, University of Rochester, Harkness Hall, Rochester, NY 14627, USA. Tel.: +1 585 275 1871.
E-mail address: yongsung.chang@gmail.com (Y. Chang).

0304-3932/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.jmoneco.2012.01.001

so as to capture the dispersion in wage rates and hours observed across workers in the SIPP data. The size of shocks to match-specific productivity and the cost of posting a vacancy are calibrated so that the model's steady state exactly matches the average rates of separation, job finding, and employment for each group in the data.

Section 5 examines the model's cyclical predictions for our four groups. The model predicts that workers with weak comparative advantage, that is, low-hours and wages, will exhibit much more cyclical separations, finding rates, and employment. Compared to the model, the data show that separation rates are much more cyclical for high-wage workers and that finding rates are much more cyclical for high-hours workers. The model does most poorly in describing the cyclicality of finding rates across low- versus high-hours workers. The model predicts that low-hours workers will contribute most of the cyclicality in the aggregate finding rate; but for the data workers who work above median hours actually contribute more than half of the cyclicality in the aggregate finding rate.

Our model produces volatility in aggregate unemployment that is 75% higher than the model without heterogeneity with the same average replacement rate. This reflects the high employment volatility of low-wage, low-hours workers that is not offset by the lower volatility for workers with higher wages and hours. But the model still substantially underpredicts the volatility of employment seen in the aggregated data. The analysis is also extended to consider cyclicality for women. The model continues to poorly predict aggregate cyclicality when calibrated to groups of women as well as men.

In sum, the model does not generate observed cyclicality of turnover and employment, failing most notably for higher wage and higher hours workers. Section 6 discusses the robustness of these results to our calibration choices. The last section suggests ways to extend worker heterogeneity to further test the DMP model against cross-sectional data.

2. Model

Unemployment transitions are modeled with endogenous separations and vacancy creation, as in Mortensen and Pissarides (1994), but allowing for heterogeneity in workers' market skills and values of non-market time. An intensive margin of labor supply is also introduced; this margin is exploited to match heterogeneity in labor supply for the model to that in the SIPP data.

2.1. Environment

There is a continuum of infinitely lived workers. Each worker has preferences defined by

$$E_0 \sum_{t=0}^{\infty} \beta^t \{c_{mt} + c_{nt}\},$$

where c_{mt} and c_{nt} are, respectively, consumption of a traded, market-produced good and a non-traded home-produced good. Consumption of the non-traded, home-produced, good is introduced in order to incorporate labor supply heterogeneity into the model. Following Mortensen and Pissarides, utility is linear in consumption. The time discount factor is denoted by β. It is assumed that the market equates $(1/(1+r))$, where r is the rate of return on consumption loans, to this discount factor; so consumers are indifferent to consuming or saving their wage earnings.

Workers differ in terms of ability in the market and at home. The market ability is denoted by a. A worker's productivity at home is given by ab. So *relative* productivity at home is b. A worker with a low value for b will have comparative advantage in the market (i.e., high rents to market work). The model will map high-wage workers to high values of a and high-hours workers to low values for b.

The value of home activity reflects time not spent in market work according to

$$c_{nt} = ab \cdot \frac{(1-h_t)^{1-1/\gamma} - 1}{1 - \dfrac{1}{\gamma}},$$

where h_t are market hours and $\gamma(>0)$ is finite, implying diminishing returns to non-market time, $1-h_t$, for the home activity.

There is also a continuum of identical agents who are entrepreneurs (or firms). Entrepreneurs have the ability to create job vacancies with a cost κ per vacancy. Entrepreneurs maximize the discounted value of profits

$$E_0 \sum_{t=0}^{\infty} \beta^t \pi_t.$$

A worker, when working, earns wages w_t. Note that w_t refers to the wage payment per period of employment, not the rate per hour. The hourly wage rate is w_t/h_t. This wage will differ across the workers, reflecting differences in a, b, and match quality as discussed below. These earnings are used to consume market goods. Some expenditures, however, are required to be employed, e.g., for transportation or clothing, that are not valued in c_{mt}. These expenditures equal ω per period employed. Because they constitute a smaller share of earnings for high-wage, high-hours workers, this is an added source of market comparative advantage for these workers. If unemployed, a worker receives an unemployment benefit of ϕ which is assumed to be proportional to a worker's average earnings.

A matched pair produces output, $y_t = a x_t z_t h_t$, where a is the worker's ability, x_t is idiosyncratic match-specific productivity (i.e., match quality), z_t is aggregate productivity, and h_t are market hours worked. Idiosyncratic match productivity and aggregate productivity evolve over time according to Markov processes, respectively $\Pr[x_{t+1} < x' | x_t = x] = F(x'|x)$ and $\Pr[z_{t+1} < z' | z_t = z] = D(z'|z)$.

Matching markets are segmented by worker type (a,b). These separate markets can be interpreted as search and matching directed by skill and by desired hours, as workers with a high value of home time will be interested in shorter hours worked. The vacancy cost κ is allowed to differ across labor markets; so $\kappa = \kappa(a,b)$. Guerrieri et al. (2009) characterize separating contracts in a search environment, such as here, with distinct types. For our model, as calibrated below, workers in fact do not want to "cross" the market—i.e., high-b (low-labor-supply) workers have higher expected utility in their prescribed market than they would expect in the low-b market, working high-hours.

The number of new meetings between the unemployed and vacancies in each market is determined by a matching function, $m_{it} = \eta u_{it}^{1-\alpha} v_{it}^{\alpha}$, where v is the number of vacancies, while u is the number of unemployed workers. Each market is indexed by i where i reflects the worker type (combinations of a and b). The matching rate for an unemployed worker is $p(\theta_t) = m_t/u_t = \eta\theta_t^{\alpha}$, where $\theta_t = v_t/u_t$ is the vacancy–unemployment ratio, i.e., labor market tightness, with index i implicit. The probability that a vacant job matches with a worker is $q(\theta_t) = m_t/v_t = \eta\theta_t^{\alpha-1}$.

Match surplus reflects the value of the match relative to the sum of the worker's value of being unemployed and the entrepreneur's value of an unmatched vacancy (which is zero in equilibrium). There are no bargaining rigidities. Therefore, separations are efficient—occurring if and only if match surplus falls below zero. Furthermore, the choice of hours worked within the match is efficient, maximizing the match value.

The timing of events is as follows. (1) Each period matches from the previous period's search and matching are realized. Also aggregate productivity z and each match's idiosyncratic productivity x are realized. (2) Upon observing x and z, workers and entrepreneurs decide whether to continue as an employed match. Workers breaking up with an entrepreneur become unemployed with the match permanently ended. (3) For matched workers hours and wages are chosen and production takes place. Hours are chosen to maximize match surplus with the wage reflecting worker-firm bargaining. Concurrent with production, unemployed workers and vacant firms engage in the search/matching process.

2.2. Value functions and choices for hours, separations, and wages

The assumptions above of linear utility in consumption, linear production in labor, and a constant returns to scale matching function imply that choices for vacancies, separations, hours, and wages in the market for one labor group are independent of choices and outcomes in the other labor markets. For simplicity the market index i is omitted for this section. Time subscripts are also omitted: variables are understood to refer to time period t, unless marked with a prime ($'$) denoting period $t+1$.

First, consider the choice of hours. Firms and workers bargain efficiently, maximizing the value of match surplus. This requires choosing hours to equate the marginal product of an hour in the market to its marginal benefit at home: $axz = ab(1-h)^{-1/\gamma}$. So optimal hours at the intensive margin for a worker are $h^* = 1-(b/xz)^{\gamma}$. Our specification yields a Frisch elasticity of labor supply for market hours h_t of $\gamma((1-h_t)/h_t)$.

Turning to the value functions, a worker's valuation of being employed is

$$W(x,z) = (w(x,z)-\omega) + \frac{ab}{1-1/\gamma}\left[\left(\frac{b}{xz}\right)^{\gamma-1} - 1\right] + \beta E[\max\{W(x',z'),U(z')\}|x,z].$$

The expenditures necessitated by employment, ω, are netted from the wage payment. The value of home production reflects the optimal choice of market hours h^*. The maximization problem implicit in $W(x,z)$ is to choose a cut-off value, x^*, such that the match persists only if match quality x exceeds that value.

The value of being unemployed is

$$U(z) = \phi(a,b) + \beta(1-p(\theta))E[U(z')|z] + \beta p(\theta)E[W(\overline{x},z')|z].$$

Note that the functional form of c_{nt} implicitly normalizes home production to zero for the unemployed. New matches are assumed to begin with a match quality equal to the mean, $\overline{x}$, for the unconditional distribution of x. For the range of parameter values considered, this ensures that workers will in fact accept new matches.

For an entrepreneur the value of a matched job is

$$J(x,z) = axz\left[1-\left(\frac{b}{xz}\right)^{\gamma}\right] - w(x,z) + \beta E[\max\{J(x',z'),V(z')\}|x,z].$$

The value for current production reflects the optimal choice for hours. The value of a matched job J reflects the option value of being able to end the match for $t+1$ if match quality falls below x^*.

The value of a vacancy is

$$V(z) = -\kappa(a,b) + \beta q(\theta)E[J(\overline{x},z')|z] + \beta(1-q(\theta))E[V(z')|z],$$

where κ is the vacancy posting cost and $q(\theta)$ is the probability that a vacancy is filled. $V(z)$ will equal zero in equilibrium given free entry in creating vacancies.

The wage payment is assumed to reflect Nash bargaining over match surplus

$$\underset{w}{\operatorname{argmax}}(W(x,z)-U(z))^{\chi}(J(x,z)-V(z))^{1-\chi},$$

where $0 \leq \chi \leq 1$ reflects the worker's share of match surplus. This wage payment is predictably increasing in ability a, relative home productivity, b, match productivity x, and aggregate productivity z. The impacts of x and z on the wage payment reflect not only their direct roles in productivity, but also their positive impacts on hours worked.

3. Cross-sectional patterns from SIPP data

The SIPP data are described, then used to construct statistics on employment and turnover for four distinct groups based on workers' average wages and hours worked.

3.1. Our SIPP sample

The SIPP is a survey of households designed to be representative of the U.S. population. It is a series of overlapping panels, each about three years in duration, with about 20,000 households. Households are interviewed every four months. At each interview, information on work experience is collected for the three preceding as well as most recent month. The first panel was initiated in October 1983. Each year through 1993 a new panel was begun. New, slightly longer panels were initiated in 1996 and again in 2001. Our analysis pools 11 panels, excluding the panel for 1989, which is very short in duration. For our purposes the SIPP has some distinct advantages. Compared to the CPS, its panel structure allows us to construct workers' long-term average wages and hours.[1] It has a larger and more representative sample than the PSID or NLS panels.

The benchmark sample is restricted to men, though results for women are discussed as well. A cost of excluding women is that it reduces heterogeneity—women on average work fewer hours and have lower earnings than men. On this basis, the model would anticipate greater cyclicality of employment for women. In fact, it is well known that women exhibit less cyclical employment. The contrast in cyclicality between men and women is complicated, however, by the fact that men are employed at a much higher rate in cyclically sensitive industries. Shin (2000) calculates that industry composition accounts for more than 100% of the greater cyclicality for men; that is, within-industry employment is more cyclical for women. Similarly, Sahin et al. (2010) show that although men experienced a much bigger increase in unemployment during the recession that began in December 2007, this difference largely reflects the much greater share of men's employment in sectors that most declined during the recession.

Our sample is restricted to men ages 20–60. Individuals must not be in the armed forces, not disabled, not be attending school full-time, and must have remained in the survey for at least a year. The analysis is further restricted to those who averaged at least one month of employment per year and who have data on both hours worked and earnings for at least one month.[2] The resulting sample for men consists of 73,427 separate individuals, representing data on employment status for 2,136,614 monthly observations. The sample for women reflects 80,930 individuals and 2,430,376 monthly observations.

A worker is matched if he reports being with a job the entire month and reports no weeks primarily involved in search. A worker who is temporarily away from work is also treated as matched provided he returns to the same employer within three months and reports no weeks of searching. In this case, weeks not actively working are reflected in the worker's measured hours worked conditional on being matched—the intensive margin. Workers not employed in a match are referred to below as unemployed. The sample of men averages an employment rate of 92.8%. A separation corresponds to moving from being matched to being unmatched; a job finding is a transition from unmatched to matched. These rates average, respectively, 1.5% and 19.4% monthly for the sample of men. For women, the employment, separation, and finding rates average 86.4%, 1.9%, and 12.4%.

Hours worked (the intensive margin) also capture occurrences of temporary layoffs or other stretches temporarily away from a job. But variability in this component, temporary layoffs, contributes little to the variability of the intensive margin. The wage rate is measured by the hourly rate of pay on the main job. More than 60% of workers report a wage in this form. For the others, an hourly rate is constructed from their reported hours and earnings, based on how the hourly wage projects on these variables for those who do report an hourly wage.

3.2. Employment and turnover by average wages and hours in the SIPP

For each worker (ln) hourly wage rates and (ln) hours worked are averaged over all months employed.[3] Workers are sorted into one of four bins based on whether their average wage is above or below the median value and whether their

[1] Fujita et al. (2007) present detailed results on the cyclicality of separation and finding rates for workers in the SIPP, comparing these to patterns in the CPS.

[2] Self-employed workers are treated as employed, but wage rates and hours worked are based only on months working for an employer with usual weekly hours of at least 10. The SIPP interviews provide distinct answers on employment status for each month. But data for wage rates and hours only capture the interview month. Therefore, attention is restricted to the survey month in examining the cyclicality of hours.

[3] The worker's hourly wage and hours worked are projected on time dummies to obtain the worker's wage, hours, and layoff variable relative to other's for that month. The residuals are then averaged across months for an individual to obtain his average wage and hours worked.

Please cite this article as: Bils, M., et al., Comparative advantage and unemployment. Journal of Monetary Economics (2012), doi:10.1016/j.jmoneco.2012.01.001

Table 1
Sample shares and deviations from sample means in long-term wages and hours by group, SIPP data.

Hours group	Wage group	
	Low wage (%)	High wage (%)
Low hours	Share: 26.7 Wage: −37 Hours: −15	23.3 +34 −10
High hours	23.3 −30 +13	26.7 +34 +12

Sample based on 73,427 men. Overall means are 2.70 for ln(wage) and 5.23 for ln(hours). Overall standard deviations are 0.42 for ln(wage) and 0.19 for ln(hours). Correlation between long-term wage and hours equals 0.15.

Table 2
Employment, separation, and finding rates by group, SIPP data.

Hours group	Wage group	
	Low wage (%)	High wage (%)
Low hours	Employment: 87.5 Separations: 2.46 Findings: 17.4	94.8 1.02 16.5
High hours	91.9 2.02 23.2	96.6 0.84 22.6

Overall means are 92.8% for employment (7.2% for non-employment rate), 1.54% for separation rate, and 19.4% for finding rate.

average hours are above or below the median. The median wage for men is $15.26 per hour in January 2009 dollars. The median hours worked is 180 per month. The standard deviations equal 42% for the hourly wage and 19% for hours worked. The correlation between average wage and average hours is positive, but fairly small, at 0.15. Statistics for the four groups of men, low-wage/low-hours, low-wage/high-hours, and so forth, broken at the medians of wages and hours are contained in Tables 1 and 2. For women, the same statistics, due to page restraints, are given in appendix tables available on the JME web site.[4] But we discuss these statistics below. Groups will often be referred to by their location in the tables, e.g., northwest group for the low-wage and low-hours workers, southeast for high-wage, high-hours.

Table 1 first reports each group's share. Reflecting the modest positive correlation between average wage and hours, the diagonal groups, low-wage/low-hours and high-wage/high-hours, are modestly larger, each at 27% of the sample, than the off-diagonal groups, each at 23%. But the off-diagonal groups still combine for nearly half of the sample. For this reason, it would greatly misrepresent the data to model heterogeneity in labor supply as captured only by heterogeneity in market skills. Furthermore, it is shown below that breaking the sample by hours worked is important for capturing differences in cyclical turnover (separation and finding rates). For women, 30% of the sample is represented in each of the groups along the diagonal, with 20% in each of the off-diagonal groups.

Table 1 also reports for men each group's mean average (ln) wage and mean average (ln) hours, both expressed as their deviation from the mean for the entire sample. Overall, the high-wage men exhibit 68% higher wages than the low-wage men. High-hours men, overall, work 25% more market hours than low-hours men. For women, high-wage workers average 65% higher wages than low-wage; high-hours female workers average 41% higher hours than low-hours women.

The appendix reports distributions of schooling attainment and age for each of the four groups as well as shares of men and women in cyclical industries. It is worth noting that the share of men employed in construction or durable manufacturing (cyclical industries) is similar across the four groups. So industry composition presumably explains little of the differences in cyclicality of employment across the four groups. For women, the share of workers in cyclical industries is much lower, at 7%, than the 24% for men. Furthermore, it is considerably higher for high-hours women, at 10%, than for low-hours women, at 4%. For these reasons, comparisons of cyclicality across low- versus high-hours women may reflect the greater volatility of construction and durables manufacturing. And, as discussed above, it is particularly difficult to draw conclusions from the employment cyclicality of women relative to that of men.

[4] http://jme.rochester.edu/JMEsupmat.htm.

Please cite this article as: Bils, M., et al., Comparative advantage and unemployment. Journal of Monetary Economics (2012), doi:10.1016/j.jmoneco.2012.01.001

Table 2 reports employment and turnover rates across the four groups of men. Starting with the two extremes along the diagonal, the northwest group (low-wage/low-hours) exhibit an employment rate of only 87.5% compared to 96.6% for the southeast group (high-wage/high-hours). This differential appears more extreme if viewed in terms of the unemployment rates, representing a rate nearly four times as large (12.5%) for the northwest group compared to that (3.4%) for the southeast group. Given this extreme differential, it is not surprising that the low-wage/low-hours workers display both higher separation rates and lower finding rates. But this difference is much more striking for separation rates. The northwest group shows three times the separation rate as the opposite extreme group in the southeast (2.5% versus 0.8%), whereas their finding rate is only lower by about 25%.

Employment rates for the off-diagonal groups are intermediate to these two extremes. The employment rate for the southwest group (low-wage/high-hours) is 91.9%; for the northeast (high-wage/low-hours) it is 94.8%. Rates of separation project much more on workers' wage rates, whereas finding rates are better explained by differences in hours worked. Workers with low wages and high-hours actually show a slightly higher finding rate (23.2%) than high-wage/high-hours workers (22.6%). But the separation rate for these workers (2.0%) is much higher. Similarly, the finding rate for high-wage/low-hours workers is even lower than that for the extreme group with both low wages and hours. But their separation rate (1.0%) is much closer to the southeast group with high wages and hours.

Women on average shows a 7% lower employment rate than men, with much of this driven by a low employment rate for low-wage, low-hours women of 76.0%. By contrast, the rate for high-wage/high-hours women is 95.2%, nearly as high as the rate for high-wage/high-hours men. Thus, the non-employment rate for low-wage/low-hours workers, 24.0%, is five times the rate of 4.8% for high-wage/high-hours women. This difference is driven by the separation rate, which is four times higher for low-wage/low-hours women, at 3.6%, compared to 0.9% for high-wage/high-hours women. But the low-wage/low-hours women do show a lower finding rate, 11.5%, compared to the 15.0% rate for high-wage/high-hours women. Employment rates for the off-diagonal women's groups are 85.5% for low-wage/high-hours and 89.9% for high-wage/low-hours. Separation and finding rates for these groups are, respectively, 2.3% and 14.0% versus 1.4% and 11.6%.

4. Calibration and steady state

This section calibrates the model to make the dispersion in wages and hours comparable to that in the SIPP data. In addition, parameters are chosen so that average rates of employment, separation, and job finding across groups match those in the data.

4.1. Calibrating to SIPP data

The discount factor β is set to 0.9966, implying an annualized real interest rate of 4%. The Frisch elasticity, $\gamma((1-h)/h)$, reflects both the parameter γ and the level of hours worked. Market hours, evaluated at mean match quality equal $1-b^{\gamma}$. The values of b and γ are set to generate market hours equal to 0.5, with a Frisch elasticity of 1/3, for our high-hours workers. This requires $\gamma=1/3$ and $b=1/8$. The same value of γ is imposed across the four groups. This implies a larger Frisch elasticity for workers who work shorter hours, with the average elasticity across groups equal to 0.44. The bargaining share for workers is set to 0.5. The matching power α is also set to 0.5 so that Hosios (1990) condition holds.

The size of the four groups is matched to those in the SIPP data: 27% for the two groups along the diagonal and 23% for the off-diagonal groups. Market ability a for the higher productivity groups is normalized to one. An earnings ability for low-wage groups of $a=0.5$ makes the cross-sectional dispersion of wages across our four groups in the model comparable to that seen in the SIPP data. For the low-hours groups, home productivity (relative to market) equal to 0.25 is required to generate a cross-sectional dispersion in log hours that mimics the data. Table 3 presents the ability parameters (a,b), respectively, for each of the four groups.

The key outcomes targeted across groups are the average rates of separation and job finding in the SIPP data. In turn, these rates depend primarily on the replacement rate while unemployed, the size of idiosyncratic shocks to matches, and the vacancy posting cost. The size of the replacement rate reflects the unemployment benefit, ϕ, the expenditure saved by not working, ω, and the extra home production when unemployed. The unemployment benefit (ϕ) is assumed to be proportional to each group's earnings evaluated at match quality equal to one ($x=1$),[5] with the benefit set to 20% of earnings. Shimer (2005) assumes a replacement rate of 40%. But for his calibration this rate should reflect any gains with unemployment from increased leisure or home production, whereas here it is an explicit, separate component.

The expenditure necessitated by work, ω, is set to 0.05. This is 10% of average earnings for workers with high wages and hours and, at the other extreme, nearly 27% of earnings for workers with both low wages and low-hours. Aguiar and Hurst (2009) show that spending on food away from home, clothing, and transportation all reflect employment variation over the life-cycle. Their estimates support that 5% of consumer spending is driven by employment expenses. Other spending

[5] Anderson and Meyer (1997) report the level of unemployment benefits by wage decile based on the 1993 panel of the SIPP data. If the breakdowns in benefits by wage from Anderson and Meyer are viewed together with a breakdown in unemployment by wage, this suggests an elasticity of unemployment benefits with respect to wage that is close to one.

Table 3
Calibrated productivities, replacement rates, match shocks, and vacancy costs by group.

Hours group	Wage group	
	Low wage	High wage
Low hours	$a=0.5$ $b=0.25$ Replace rate $=98.4\%$ $\sigma_x=3.6\%$ $\kappa/q(\theta)=0.09$ $\frac{\kappa/q(\theta)}{ah}=0.46$	$a=1.0$ $b=0.25$ Replace rate $=84.8\%$ $\sigma_x=2.5\%$ $\kappa/q(\theta)=0.38$ $\frac{\kappa/q(\theta)}{ah}=1.02$
High hours	$a=0.5$ $b=0.125$ Replace rate $=77.5\%$ $\sigma_x=8.2\%$ $\kappa/q(\theta)=0.35$ $\frac{\kappa/q(\theta)}{ah}=1.40$	$a=1.0$ $b=0.125$ Replace rate $=67.5\%$ $\sigma_x=2.9\%$ $\kappa/q(\theta)=0.74$ $\frac{\kappa/q(\theta)}{ah}=1.48$

categories such as child care will also be affected by employment. In addition, any costs that fall on the employer that have a fixed per worker nature would act in the model precisely like the expenditures in ω.

Of course, some work-required spending may be higher for jobs with higher wages or hours. Consider the impact of making such spending completely proportional to earnings. Holding match shocks and vacancy costs unchanged, the average employment rate for low-wage workers becomes 92.5% versus 94.0% for high-wage workers. For the SIPP data, these rates are 89.7% for low-wage versus 95.7% for high-wage workers. Thus, if all fixed costs of employment are disallowed, the model cannot generate reasonable dispersion in employment. Our interpretation is that some fixed costs are plausible. But an alternative interpretation is that the model fails, not only for understanding the cyclicality of employment, but also for matching the steady-state employment differences across groups.

The replacement ratio reflects the flow benefits to being unemployed (home production + unemployment benefit + work-related expenditures saved) relative to the earnings of employed.[6] The gain in home production while unemployed for the two high-hours groups equals 37.5% of market productivity; for the two groups with low-hours it is 51.3%. Combined with the unemployment benefits and expenditures required for working, this yields a steady-state replacement ratio equivalent to 67.5% of market output for the southeast group. This value is considerably higher than the 40% employed by Shimer's calibration, but close to the replacement rates assumed by Costain and Reiter (2008) and by Hall (2005). Those authors, however, employ that ratio for all workers, whereas we employ it only for high-wage/high-hours workers. The replacement rates for each of the four groups is presented in Table 3. The average rate across the groups is 82%.

Shocks to match-specific productivity are assumed to be highly persistent, with autocorrelation equal to 0.98. This high persistence accords well with the high persistence typically estimated for individual wages and earnings (e.g., Topel and Ward, 1992). The volatility of these shocks, σ_x, together with the vacancy posting cost, κ, are then set in order to mimic the monthly separation and finding rates for each group in the data. For instance, for the high-wage, high-hours group this is achieved by $\sigma_x=2.7\%$ and $\kappa=0.31$. With $\sigma_x=2.7\%$ the model generates wage dispersion of 5% due to differences in match quality for the workers with strong comparative advantage (the southeast group). It generates a cross-sectional dispersion in monthly wage growth for these workers of 3%. These values are arguably conservative relative to the empirical literature.[7] Table 3 reports the calibrated values of σ_x for each of the four groups. For the two groups with low-hours (northwest and northeast) σ_x is around 3%, close to that for the southeast group. But for the southwest group (those with low wages but high-hours) the required σ_x to match the average separation rate is considerably higher at 7.6%.

For each group we report the expected vacancy posting cost per hire, $\kappa/q(\theta)$, as well as this cost expressed in terms of months of worker output, $(\kappa/q(\theta))/ah$. From Table 3, for high-hours workers this expected cost is about 1.5 months of output. For the northwest group (low-hours/low-wage) the expected cost is only about a half month's earnings. This lower cost is required for the model to explain why these workers, while having far below average earnings, exhibit a finding rate that is nearly 90% that of the overall average.

Parameters σ_x and κ are chosen to match each group's employment and turnover rates. This allows us to examine cyclicality for a benchmark model that accurately depicts employment and turnover rates across workers. But this does require higher match quality shocks for the southwest group and lower vacancy costs for the northwest group. These

[6] Evaluated at the unconditional mean for match quality ($x=1$), this ratio is $(\omega+\phi(a,b)+\gamma ab(b^{\gamma-1}-1)/(1-\gamma))/a(1-b^{\gamma})$.

[7] For instance, Topel and Ward (1992), based on administrative data, report a standard deviation in annual rates of growth in earnings of 19%. This dispersion is considerably greater than that generated by our model with $\sigma_x=2.7\%$. Increasing σ_x would cause match rents to become more disperse, making employment even less responsive to aggregate shocks.

Table 4
Deviations from means in long-term wages and hours by group, model results.

Hours group	Wage group	
	Low wage (%)	High wage (%)
Low hours	Wage: −34 Hours: −14	+34 −16
High hours	−29 +16	+29 +13

Overall standard deviations are 0.32 for ln(wage) and 0.15 for ln(hours). Correlation between long-term wage and hours equals 0.09.

Table 5
Employment, separation, and finding rates for benchmark and restricted models (in %).

Hours group	Wage group				
	Low wage			High wage	
		Benchmark	Restricted	Benchmark	Restricted
Low hours	Employment	87.7	86.1	94.1	93.3
	Separations	2.45	2.16	1.04	1.56
	Findings	17.4	13.4	16.6	21.5
High hours		92.0	94.8	96.4	96.0
		2.04	1.38	0.84	1.22
		23.3	24.9	22.6	29.0

For the benchmark model the size of match shocks, σ_x, and vacancy costs, κ, are calibrated to each group as in Table 7; for the restricted model $\sigma_x = 3.7\%$ and $(\kappa/q(\theta))/ah = 1.23$ for all groups.

parameter differences are difficult to justify a priori. For this reason, we also consider a more parsimonious version that restricts σ_x and the hiring cost relative to output to be the same for all workers. For this restrictive case, $\sigma_x = 3.7\%$ and $(\kappa/q(\theta))/ah = 1.23$. These values are required for this restricted model to match the aggregate separation, finding, and employment rates from the data.

4.2. Steady state

Table 4 presents the model's steady-state mean wages and hours worked for each of the four labor groups. Across low-hours workers, the high-wage group (northeast) displays a wage that is 68% higher than the low-wage group (northwest). Across high-hours workers the wage differential between wage groups is smaller at 58%. The difference in wages across groups in Table 4 matches fairly closely the average differential by group observed in the SIPP data. But the cross-sectional wage dispersion in the SIPP, with a standard deviation of 42%, is greater than the 32% for our model economy. Turning to hours, each high-hours group for the model displays hours worked that are about 30% higher than those for the low-hours group with comparable wage rates. This is close to the percentage found in the SIPP data. The model generates an overall standard deviation in hours worked of 15%, lower than that of 19% for the SIPP data. The cross-sectional correlation of log hours and log wages in our model is 0.09, somewhat lower than that in the SIPP (0.15). The higher standard deviations of wages and hours, and their higher correlation, in the SIPP data reflects that the data have heterogeneity in wages and hours within each of the four groups.

Table 5 presents the employment and turnover rates for our benchmark model as well as for the restricted model with common-sized match shocks and hiring costs across groups. By construction, the benchmark model matches each group's rates as given in Table 2. But the restricted model generates similar dispersion in employment rates across the four groups. Thus, most of the differences in employment is captured just by the differences in the values of market and non-market productivity by group, parameters a and b, and differences in comparative advantage driven by the fixed employment cost ω. The restricted model understates separations for low-wage workers and understates the finding rate for low-wage/low-hours workers.

The differences in separation and finding rates across the four groups are directly related to the rents to employment. These rents can be represented by the difference between the match quality, x, that a worker has in employment versus the critical match quality, x^*, at which the match would be dissolved—a match with $x = x^*$ would have zero rents. Fig. 1 presents the distributions of match rents (measured by $x - x^*$) separately for each of the groups for the benchmark model. The low-hours groups show distinctly less surplus than the high-hours groups. In particular, the fraction of matches with

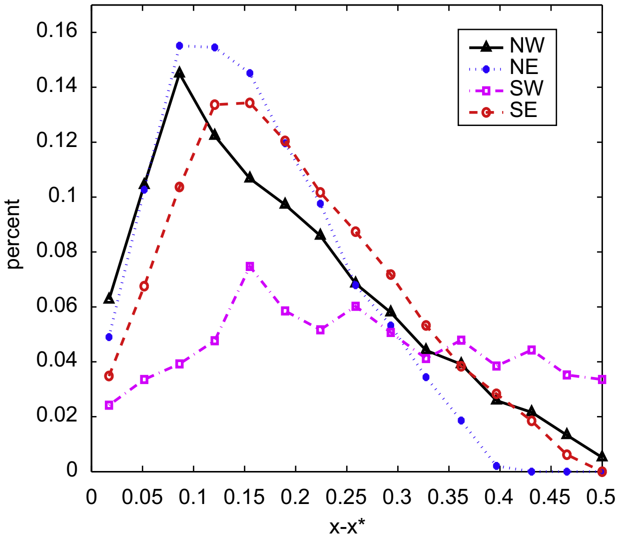

Fig. 1. Distribution of match rents. *Notes*: Distribution of match rents for each group based on the benchmark steady state.

$x-x^*$ less than 0.02 is nearly twice as large for the low-hours workers. These distributions of match rents, especially at low values of $x-x^*$, are telling for the cyclical behavior of separations and employment across the four groups examined next. Consider a 2% decrease in aggregate productivity while, for this thought experiment, holding the value of unemployment constant. This decrease would cause more than 5% of current matches to dissolve for the northwest group but only 2% of matches for the southwest group. That is, the large fraction of workers with low match rent for the low-hours groups predicts a sharp increase in the separations during a downturn. A decline in aggregate productivity also creates a greater percentage drop in surplus ($x-x^*$) for workers with low-hours; in turn the model predicts a larger reduction in vacancy creation rates and finding rates for these workers.

5. Business cycles

For each of the four groups, stratified by wage and hours, model predictions are compared to the data with respect to the (1) relative size of employment fluctuations, (2) importance of the hours versus employment fluctuations, and (3) cyclicality of separation and finding rates. This is followed by examining implications from aggregating these disparate groups, asking, in particular, to what extent the cyclical volatility of aggregate statistics is influenced by the cyclical workers with weak comparative advantage.

5.1. Fluctuations in employment and hours across workers

Business cycles for the model are created by hitting each group with persistent shocks to aggregate productivity z. We calibrate these shocks to display an autocorrelation of 0.97, with an innovation standard deviation of 0.77%, in order to mimic cyclical properties of aggregate labor productivity in the U.S. For both the SIPP and the model, time series are constructed for each of the four groups. Each time series, both from the SIPP and model-generated, is time-averaged to quarterly then HP filtered with a smoothing parameter of 1600. Quarterly series are examined partly for comparability to the literature, but also to mitigate any impact from sampling error in the SIPP data. Seasonals are also removed for the SIPP-based series.

Table 6 provides the cyclical response of employment for each group to aggregate hours. Aggregate hours equals the sum of hours across all workers; so it reflects both the intensive and extensive margins. All variables are in natural logs, i.e., the table reports the estimated coefficient from regressing the group's (ln) employment on (ln) aggregate hours. The table gives results for fluctuations for men from the SIPP data and for model-simulated fluctuations. The model statistics are presented for two versions: our benchmark model and the restricted version with common-sized match shocks and vacancy costs across groups.

First, consider the two diagonal groups. The benchmark model generates 10 times greater cyclicality in employment for the group with weak comparative advantage (low-wage/low-hours) than for those with strong comparative advantage (high-wage/high-hours), with a response of 1.1% for the first group versus 0.1% for the latter. (For the restricted model this ratio is even higher, at 14, due to even greater volatility for low-wage/low-hours workers.) This highlights our motivation for splitting the sample by hours and earnings—the DMP model suggests that the replacement rate is key to the size of fluctuations, and this rate should be considerably greater for those with lower hours and earnings. Employment is more

Table 6
Business cycles by group, data compared to models.

Hours group	Wage group	Low wage			High wage		
		SIPP data	Benchmark model	Restricted model	SIPP	Benchmark	Restricted
Low hours	Employment	1.03 (0.12)	1.09 (0.16)	1.42 (0.19)	0.57 (0.11)	0.33 (0.02)	0.31 (0.004)
	Hours	0.32 (0.07)	0.76 (0.05)	0.68 (0.05)	0.21 (0.03)	0.84 (0.04)	0.73 (0.05)
	Sep's	−4.56 (1.92)	−3.07 (0.23)	−3.94 (0.27)	−6.85 (1.98)	−2.73 (0.25)	−1.70 (0.15)
	Find's	3.19 (0.92)	7.27 (0.33)	6.45 (0.20)	4.31 (2.69)	3.63 (0.10)	2.97 (0.12)
High hours		0.45 (0.07)	0.18 (0.01)	0.18 (0.01)	0.37 (0.07)	0.11 (0.01)	0.10 (0.004)
		0.45 (0.10)	0.47 (0.03)	0.44 (0.03)	0.34 (0.05)	0.51 (0.03)	0.45 (0.03)
		−2.28 (1.50)	−0.51 (0.14)	−1.19 (0.17)	−5.79 (2.16)	−1.16 (0.18)	−0.79 (0.10)
		5.80 (1.26)	2.02 (0.09)	2.29 (0.11)	5.34 (1.85)	1.97 (0.09)	1.74 (0.10)

Responses of hours, employment, separation, and finding rates to aggregate hours (employment times hours), all in natural logs. Series are quarterly, HP-filtered with parameter 1600. SIPP data are seasonally adjusted. There are 73 observations per group. Standard errors (Newey-West corrected) are in parentheses. Model statistics are means across 100 simulations; standard deviations for the simulations are in parentheses.

cyclical for the low-wage/low-hours group in the SIPP data. But the response in employment is only two to three times as large as that for the high-wage/high-hours group.

Turning to the off-diagonal groups, the model (whether benchmark or restricted) suggests that these groups will display employment fluctuations intermediate to the two extreme groups, but closer in magnitude to that of the high-wage/high-hours workers. Both these predictions hold true qualitatively in the SIPP data. But, as with comparing the extreme groups, the magnitude of fluctuations is much more comparable across groups in the data than predicted by the model. For instance, the model predicts employment fluctuations for low-wage/high-hours workers that are less than one-sixth the magnitude for low-wage/low-hours, but in the data they are nearly half as large.

Table 6 also reports the cyclical response of hours worked to aggregate hours. It is a stylized feature of U.S. business cycles that employment fluctuations are larger in magnitude than fluctuations in the intensive hours margin. This holds for the SIPP data when aggregated as well. But Table 6 shows that this feature actually only holds for low-hours workers. For low-hours workers employment is about three times as cyclical as hours, whereas for high-hours workers the two margins are nearly equally cyclical. By contrast, the model predicts most of the labor response is in hours for every group except the low-wage, low-hours workers. And even for this group the model yields an employment response equal to about 1.5 that in hours, whereas, in the data, employment responds by more than three times as much. For each of the other three groups the model substantially underpredicts the relative cyclicality of the employment margin (again, either for the benchmark or restricted versions).

Thus, despite the model being calibrated to have high replacement rates, averaging about 80% across workers, it fails to generate the volatility of the employment margin seen in the data. While this echoes criticisms by Shimer (2005), among others, note that our conclusion is based on the *relative* responses to aggregate hours. So it should not hinge on assuming productivity shocks as the specific source of business cycle fluctuations or, even knowing the magnitudes of cyclical shocks.

5.2. *Cyclicality in separation and finding rates*

The model predicts greater cyclicality in both separation and finding rates for low-hours, low-earnings workers, especially with respect to finding rates. But the data do not show this—the separation rate is more cyclical for workers with higher wages and the finding rate is considerably more cyclical for workers who work longer hours.

Table 6, third and fourth rows of each cell, reports the responses of separation and finding rates to aggregate hours. Note that these are estimated responses in the natural logs of separation and finding rates. Looking first at separations, the benchmark model predicts that a 1% increase in aggregate hours is associated with a decrease in the separation rate of 3.1% for low-wage/low-hours workers (northwest group) and of 2.7% for high-wage/low-hours (northeast). The model predicts much less cyclical separations for the high-hours groups. In the data, separation rates for the low-hours groups are even more cyclical than predicted, with decreases of 4.6% for the northwest group and 6.9% for the northeast group. More problematic, in the data high-wage, high-hours workers also show very cyclical separations. For this group a 1% increase in aggregate hours is associated with a decrease in the separation rate of 5.8%. That elasticity response is as large as for the low-hours workers, albeit with respect to a much smaller average separation rate. The group that shows the least cyclical separations is the southwest group (low-wage/high-hours). But the benchmark model actually anticipates this; it predicts even less cyclical separations for this group, based on the large calibrated match shocks and match rents for these workers. Contrasting the two higher wage groups with the two lower wage groups, we see that the separation rate is actually more cyclical for higher wage workers. The elasticity of the separation rate with respect to aggregate hours is 6.3 for high-wage

Please cite this article as: Bils, M., et al., Comparative advantage and unemployment. Journal of Monetary Economics (2012), doi:10.1016/j.jmoneco.2012.01.001

Table 7
Average versus business cycle shares in turnover by group, data compared to model.

Hours group	Wage group	Low wage SIPP	Low wage Model	High wage SIPP	High wage Model	All wages SIPP	All wages Model
Low hours	Average share of turnover	0.40		0.15		0.55	
	Share of cyclical separations	0.41	0.62	0.24	0.22	0.65	0.84
	Share of cyclical findings	0.29	0.68	0.14	0.12	0.43	0.80
High hours		0.30		0.15		0.45	
		0.15	0.08	0.20	0.09	0.35	0.17
		0.40	0.14	0.17	0.06	0.57	0.20
All hours		0.70		0.30			
		0.56	0.70	0.44	0.31		
		0.69	0.82	0.31	0.18		

Share of cyclical separations is the response in separations to aggregate total hours in that group, relative to the total response in separations across all four groups. The share of cyclical findings is defined comparably.

workers compared to 3.5 for low-wage, with much of this driven by the lower cyclicality of separations for low-wage/high-hours workers. Mueller (2010) examines the cyclicality of separations in the CPS by wage rates and finds, consistent with our results, that separations are more cyclical for men with higher wages.

Turning to finding rates, both the data and model display procyclical finding rates for all four groups; but its importance across groups differs greatly between the model and the data. For the most cyclical workers (low-wage/low-hours group) the model predicts that the finding rate increases by 7.3% for a 1% increase in aggregate hours, whereas this rate increases by only 3.2% in the data. At the same time, the model fails to predict the extreme cyclicality in finding rates for high-hours workers. Both groups of high-hours workers exhibit a cyclical elasticity of the finding rate that is three times that predicted by the model.

To summarize Table 6, the model does not generate enough cyclicality in separations for high-wage workers and does not generate enough cyclicality in finding rates for high-hours workers. Thus, it fails to generate observed cyclicality in separations or findings for workers with both high-wage rates and high-hours, that is, those workers with the strongest comparative advantage in the market.

In Table 7 we conduct a business cycle accounting of the contributions of the four groups to cyclicality in overall separation and finding rates. Table 7 shows three rows of numbers for each of the four groups, with each row having one entry for the data and one for the benchmark model. The first row gives the average share of each group in turnover, separations or findings. The second row gives the share of each group in the predicted decrease in separations that occurs with a cyclical expansion (an increase in aggregate hours). These shares equal the cyclical elasticity of the separation rate for each group from Table 6 weighted by the long-run share of that group in separations divided by the sum of these weighted elasticities across all groups. For instance, for low-wage, low-hours workers the contribution to cyclicality in separations in the SIPP data, 0.41, equals $(-4.56*0.40)/((-4.56*0.40)+(-6.85*0.15)+(-0.51*0.30)+(-5.79*0.15))$. The third row gives the share of each group in the predicted increase in findings that occurs with a cyclical boom. It is calculated using the elasticities for findings from Table 6 in a manner that parallels that just described for separations. Also, columns at the far bottom and far right of the table provide the same statistics aggregating the groups according to wage group (far bottom) or hours group (far right).

First, consider the model predictions. The model predicts that about two-thirds of the cyclicality in turnover rates is contributed just by the low-wage, low-hours group. This reflects the model's prediction of more cyclically elastic turnover rates for this group as well as its high share, 40%, in average turnover. The model predicts little contribution from high-wage/high-hours workers, with a share of 0.09 for separations and 0.06 for findings. Combining the two low-hours groups, the model predicts that these groups contribute 84% and 80%, respectively, of the cyclicality in separation and finding rates. The two high-hours groups contribute little: 16% of cyclicality for separations and 20% for findings.

The data give a much different picture. From Table 6, the high-wage/high-hours group exhibits both more cyclical separations and findings than predicted. As a result, their contributions to the overall cyclicality of the turnover rates are more than twice that predicted by the model. It does remain true, however, that the low-wage/low-hours group contributes about twice as much to the cyclicality of separation and finding rates than do the high-wage/high-hours, thanks to the large share of this group in average turnover. Similarly, although separation rates are more cyclical for the high-wage groups than for the low-wage groups, the low-wage groups still contribute more (56%) to the overall cyclicality of separations, reflecting their predominance in separations. But if the four groups are split into the two high-hours versus the two low-hours groups, the high-hours groups actually contribute more (57%) to the cyclicality of the overall finding rate than low-hours groups (43%). We see this as the most striking contradiction of the model, as the model predicts that the high-hours groups contribute only 20% of cyclicality in findings.

Please cite this article as: Bils, M., et al., Comparative advantage and unemployment. Journal of Monetary Economics (2012), doi:10.1016/j.jmoneco.2012.01.001

Table 8
Aggregate business cycles, data compared to model.

Variable	SIPP data	Benchmark model	Restricted model
Employment	0.66 (0.05)	0.42 (0.04)	0.50 (0.04)
Hours	0.34 (0.05)	0.58 (0.04)	0.50 (0.04)
Separation rate	−4.08 (1.16)	−1.85 (0.15)	−2.05 (0.11)
Finding rate	4.57 (0.87)	4.68 (0.19)	4.60 (0.12)

Responses of separation rate and finding rate to aggregate total hours (employment times hours), all in natural logs. Series are quarterly, HP-filtered with parameter 1600. SIPP data are seasonally adjusted. There are 73 observations per group. Standard errors (Newey-West corrected) are in parentheses. Model statistics are means across 100 simulations; standard deviations for the simulations are in parentheses.

5.3. Aggregation effects

Our focus has been comparing predictions of the DMP model for cyclicality *across* workers with that in the data. Table 8 presents the same statistics—cyclicality of employment, hours, and turnover rates—for the aggregate economy. Aggregation effectively puts more weight on the low-wage/low-hours workers due to their much greater cyclicality. Table 8 shows whether these compositional effects act to magnify or disguise the shortcomings of the model, presenting statistics aggregating groups for the SIPP data, then for both our benchmark model economy and its restricted version.

First, consider the relative magnitudes for fluctuations in employment and hours. In the data, cyclical fluctuations in employment are more important, contributing 66% of fluctuations in total hours. The model predicts that cyclical fluctuations are actually reflected more in hours (58%) than in employment (42%). Aggregating does modestly increase the importance of the employment margin for the model compared to the weighted average of its importance across the four groups, reported in Table 6, which equals 0.34. This reflects the disproportionate weight of the low-wage/low-hours workers in aggregate fluctuations. For this group the model does generate that the majority of fluctuations occur through the extensive margin. (For the restricted model this effect is even greater.)

Next, consider the responses of aggregate turnover rates to total labor hours for both the data and the model. For the benchmark model a 1% increase in aggregate hours is associated with a decrease in the separation rate of 1.9% and an increase in the finding rate of 4.7%; so the finding rate is predicted to be much more cyclical. The restricted model generates somewhat more cyclical separations than the benchmark model, but the cyclical elasticity of the finding rate remains twice that for the separation rate. The aggregated SIPP data show a very cyclical finding rate: a 1% increase in aggregate hours is associated with an increase in the finding rate of 4.6%, almost the same as for the model. But, unlike the model, the data also show a very cyclical separation rate, with the separation rate showing an elasticity of −4.1 with respect to aggregate hours.

Aggregation partially hides the failures of the model. Aggregating finding rates across the groups puts a disproportionate weight on the low-wage, low-hours groups because they make up nearly half of the unemployed; and the model predicts an extremely cyclical finding rate for this group. The weighted average of the elasticities in Table 6 is 3.8, but aggregating produces the elasticity of 4.7 in Table 8, in line with the data. For the data, aggregating reduces the elasticity of the separation rate from −4.9 (the weighted average of the elasticities of the four groups) to −4.1. But the latter number remains considerably higher than predicted by the model. Aggregating acts to reduce this elasticity because, for the data, the separation rate is more cyclical for the higher wage groups, those that exhibit much lower average separation rates.

5.4. Results for women

Appendix tables provide parallel results for women. The model is re-calibrated to fit the women's sample from the SIPP. Given lower wage rates for women, market productivities are scaled down by a factor of 0.74 relative to the calibration for men in Table 3. Non-market productivities, *b*, are 0.22 for low-hours and 0.13 for high-hours women. The calibrated replacement rates for the women's groups are higher than those for men. For instance, the calibrated replacement rate for high-wage/high-hours women is 72%, and for low-wage/low-hours women, it is 102%. As for men, match-specific shocks and vacancy costs are calibrated to match the average turnover rates for each of the four groups of women. The vacancy costs per job created are similar across groups and similar to that for men, equaling 6–8 weeks of output across groups. For high-wage/high-hours women the standard deviations for match-specific shocks is similar to that for men, with $\sigma_x = 3.5\%$. It is higher for the other groups. In particular, the high separation rates for both low-wage groups require a much larger value for σ_x, equal to 15% for each group.[8]

The SIPP data for women show that employment and hours are most cyclical for low-wage women; in fact, the cyclicality of employment is nearly all concentrated on the two low-wage groups. Each of these groups displays a cyclical

[8] The calibration cannot produce hours for the low-wage/low-hours group that are 18% below the mean for all women, yet still hit their employment rate of 76%. Reaching both targets requires such large match shocks that workers choose to keep jobs at very low match quality, working zero hours, for the opportunity that match quality will become much higher in some period: σ_x is limited to 15% so that workers only keep matches that imply positive hours.

Table 9
Cyclical statistics for model economies.

Variable	Benchmark	Restricted	No heterogeneity	Benchmark men and women	Restricted men and women
σ_u/σ_z	3.50	3.55	2.04	2.20	3.38
σ_v/σ_z	2.82	2.25	1.71	2.49	2.54
ρ_{uv}	−0.60	−0.58	−0.55	−0.59	−0.59
σ_s/σ_z	2.56	2.64	1.65	1.82	2.50
σ_f/σ_z	3.04	3.04	1.65	2.45	3.14
ρ_{sf}	−0.53	−0.52	−0.53	−0.52	−0.51

Model statistics are means across 100 simulations. Simulated series are quarterly, HP-filtered with parameter 1600. For comparison, statistics reported by Shimer (2005) for these six statistics are respectively top to bottom: 9.5, 10.1, −0.89, 3.8, 5.9, and −0.57. But note that Shimer detrends his quarterly data using HP-filter parameter 100,000.

elasticity of employment with respect to aggregate women's hours of about 0.8, compared to elasticities of only 0.1 for each of the high-wage groups. The data show little cyclicality in the finding rate for any of the groups. But the separation rate is quite counter-cyclical.

The benchmark model does predict more cyclical employment for low-wage women. But it fails to predict the magnitude of this difference. The model predicts employment that is three times as cyclical for low-wage/low-hours as for high-wage/high-hours, but in the data, it is actually eight times as cyclical. The model does not predict very cyclical separations or employment for low-wage/low-hours women despite high replacement rate. The reason for this is the high volatility of match shocks calibrated for low-wage women in order to match their high average separation rates. These shocks imply that low-wage women, conditional on being employed, have largely sorted into good matches. As a result, an aggregate shock to productivity is not likely to induce a separation.

5.5. *Unconditional volatility of unemployment produced by the model*

Shimer (2005) stresses that his calibrated DMP model generates a standard deviation for (ln)unemployment relative to that for labor productivity that is about one-half, whereas in the data this ratio is about 10. We have not emphasized this statistic. Because one cannot say that the only disturbance to employment is productivity shocks, it is difficult to judge a model by unemployment's volatility relative only to volatility in measured productivity. But it is useful to report these relative volatilities to facilitate comparing our model to others in the literature. Table 9 reports the benchmark model's standard deviation of unemployment relative to that of productivity as well as other business cycle statistics. The table also provides results for our restricted model and for a model comparably calibrated but with no heterogeneity. Last, it gives results for model economies that pool samples of men and women generated by calibrating the model sequentially for men and women.

From the first column of Table 9, the benchmark economy generates a standard deviation for unemployment, relative to that of labor productivity, of 3.5. This is much larger than for Shimer's calibration, which is about 0.5. But it is, nevertheless, only about a third that observed for the U.S. data. The volatility for our restricted economy is similar to that for the benchmark. There are three reasons why our model generates larger unemployment volatility than Shimer's. Most important, Shimer assumes a 40% replacement rate, whereas the average of this rate is about 80%. Second, Shimer assumes a constant separation rate, whereas our model's separation rate is counter-cyclical and with a standard deviation that is nearly three-quarters that for the unemployment rate. Third, because volatility increases non-linearly with the replacement rate, the very high volatility for our low-wage/low-hours group is not offset by the low volatility of the high-wage/high-hours group.

To illustrate this last point the third column of Table 9 presents results for a model economy like ours, but that is calibrated with no heterogeneity across groups. All workers are given the mean values for a and b (respectively, 0.75 and 0.1875). The replacement rate is set to 81% across all workers, comparable to the mean for the benchmark. The size of match shocks and hiring costs is chosen so that this economy mimics the same aggregate steady-state separation and finding rates (and employment rate) as the benchmark economy. With no heterogeneity the economy is much less cyclical. In particular, the standard deviation in unemployment is increased by about 75% going from the economy with no heterogeneity to our benchmark (or restricted) economies. While this result is striking, keep in mind that the mechanism driving this, the model's prediction of much greater *relative* cyclicality for low-wage/low-hours group, is little manifested in the actual data.

Turning back to the benchmark economy in column 1, the economy generates a Beveridge curve correlation of −0.60. This is well below that observed in the U.S. data, but relatively high for a DMP model with significant cyclicality in separations. Both the benchmark and restricted economies generate a separation rate that is nearly as volatile as the finding rate. These separation and finding rates are similarly negatively correlated, about −0.5, for all model economies we consider.

Please cite this article as: Bils, M., et al., Comparative advantage and unemployment. Journal of Monetary Economics (2012), doi:10.1016/j.jmoneco.2012.01.001

The model economies presented in the first three columns of Table 9 are calibrated to men. Given that nearly half of workers in the U.S. economy are women, it is natural to ask what business cycles the model generates if we pool the fluctuations predicted by the model for men with those for women. We do this in the final two columns of the table, first for our benchmark specifications, then for the more restrictive specifications that impose the constraint that men exhibit the same parameters for match shocks and hiring costs and, similarly, women share the same parameters.

Looking at column 4, pooling men's and women's labor markets actually results in smaller fluctuations in unemployment, with the standard deviation of unemployment, cut from 3.5% to 2.2%. This reflects the fact that the standard deviation of unemployment for the benchmark women's economy is only 40% as large as that for men. This may seem surprising given that the average replacement rate is higher for women. Women in the model display less volatile unemployment because they are calibrated to exhibit larger match-specific shocks than men, much larger for low-wage women. As a result, women are presumed to have typically sorted into good matches, reducing the volatility of separations and employment. The last column pools the economies for men and women generated from their restricted models. This model exhibits much lower match-specific shocks for low-wage women. As a result, it generates nearly as much volatility for unemployment as the economies calibrated only to men.

6. Robustness to alternative specifications

Several variations on the model calibration are discussed that could conceivably help to close the mismatch between the model and data. The variations considered are a skill distribution more skewed toward the weak comparative advantage group, higher replacement ratios, and varying the Frisch elasticity.

Cyclicality, both for the model and the data, is greatest for the group with the weakest comparative advantage (low-wage/low-hours workers). This raises the question of whether the model could generate more cyclical employment, better matching aggregate U.S. data, if workers are partitioned to allocate more workers to the low-wage/low-hours group. We find the answer to be no. To increase the calibrated share of low-wage/low-hours workers, it is necessary to raise the maximal wage and maximal hours that define this group. While putting more workers in the northwest group, this also produces higher average wages and hours in each of the four groups. Calibrating to these new groupings will yield greater employment rents within each group and, for this reason, generate less cyclical employment *within each group*. For instance, consider statistics constructed for four groups partitioned by whether average wages and hours are above the 75th percentile of each distribution. The newly defined northwest group (low-wage, low-hours) is of course now much larger, at 57% of the sample. But this group now displays a 19% higher average wage and 7% higher average hours. As result, when parameters a and b are calibrated to hit these higher earnings, the northwest group exhibits less than half as much cyclicality in its separation, finding, and unemployment rates. Furthermore, earnings go up considerably for each of the other three groups; so the model predicts less cyclicality for each of these groups as well.

Employment cyclicality for the model increases with the replacement rates specified for the unemployed. But note that our calibrated model already has high replacement rates: The replacement rate for workers with the strongest comparative advantage (high-wage/high-hours workers) is 67%; workers in the two off-diagonal groups average about 80% replacement rates; and workers with the weakest comparative advantage exhibit a 98% replacement rate. More important, increasing replacement rates, for example, by raising the fraction of earnings replaced by unemployment insurance, is not a reasonable solution. To generate a cyclical response in employment (relative to hours) for high-wage/high-hours workers like that in the data requires doubling unemployment insurance, from 20% to 40% of earnings. But calibrating unemployment insurance at this level, while respecting the wage and hours differences across workers, drives the total replacement rate on average above 100% for the other three groups. In particular, the replacement rate for workers with low wages and hours goes well above 100%. As a result, the model predicts zero employment for these workers, whereas in the data their employment rate is over 80%.

The Frisch elasticities for hours assumed above equal one-third for high-hours workers and a little over one-half for workers with lower hours. One might conjecture that the model's failure to capture the relative importance of the extensive versus intensive margins could be fixed by assuming smaller Frisch elasticities. But this is problematic. Reducing the Frisch elasticity makes market and non-market work poorer substitutes, which acts to increase the gains from employment. For instance, cutting the Frisch elasticities in half reduces the effective replacement value from leisure by 18% of earnings for high-hours workers and by 23% of earnings for low-hours workers. As a result, the model will generate smaller employment fluctuations for all groups; so reducing the Frisch elasticity does not correct the failure of the model to predict cyclicality of employment.

7. Conclusions

Key to generating large fluctuations in the DMP model is the presence of sufficiently many workers who have little comparative advantage in the market (low rents to employment). An economy was constructed with a realistic cross-sectional distribution of workers' comparative advantage. Workers with market comparative advantage were mapped to those earning high wages and working high-hours conditional on being employed. This allowed us to calibrate our model to the cross-sectional distribution of average wages and hours in panel data (the SIPP data).

When aggregate business cycles are introduced, the model predicts much more cyclical separations and employment for workers with weak comparative advantage, those with low wages and hours. But in the data even men who typically earn high wages and work long hours exhibit important cyclical fluctuations. The model underpredicts the cyclicality of the separation rate for high-wage workers. It underpredicts the cyclicality of the finding rate for high-hours workers: high-hours workers actually contribute more than low-hours workers to the cyclicality of the finding rate in the data. As a result, the model fails dramatically in explaining the cyclicality of employment for high-wage/high-hours workers. The model comes up short in explaining aggregate employment volatility as well. It fails to explain the greater cyclicality of the extensive employment margin than the hours margin. Although heterogeneity in the model substantially increases the volatility of employment and vacancy rates, this volatility remains much less than that observed in the U.S. data. The model fails on this score despite being calibrated with very high replacement rates for a large set of workers. Furthermore, the model's channel for increasing aggregate volatility from heterogeneity, much greater employment volatility for low-hours/low-earnings workers, is undercut by the fact that the heterogeneity in responses is much weaker in the data.

Heterogeneity in the data in wage rates and hours was exploited as a natural proxy for identifying those workers with greater rents from employment. With sufficiently rich and lengthy panel data, it would be possible to go beyond these measures. In our model, rents from employment would project on match quality (x), even conditional on a worker's wage and hours. Match quality could be partially inferred from a worker's realized tenure with an employer or, better yet, from the rate of growth in wages within the current employment match. In fact tenure, or realized wage growth on a job, should reflect match rents in more general environments than considered here, for instance, if one allows for investments in job-specific human capital or if there is learning about match quality (e.g., as in Pries, 2004). Given that the SIPP's panels are fairly short, typically only three years, it is not ideal to measure realized tenure or heterogeneity in wage profiles. But longer panels, such as the PSID or NLS data, could be exploited, though each suffers relative to the SIPP in terms of frequency of interviews and quality of turnover information.

Our analysis focused primarily on cyclicality across men. Given that men and women differ in average hours and earnings, this provides an avenue to expand on dispersion in market comparative advantage. The lower wages, hours, and employment rates for women would suggest women exhibit lower rents from employment. Based on that, women should be expected to exhibit greater cyclicality. In fact, women exhibit considerably less cyclical employment. Shin (2000) and Sahin et al. (2010) show that it is very important to correct for the cyclicality of the sector of employment, with men working in more cyclical industries. In our calibration much larger idiosyncratic shocks are required for women to explain their turnover rates. In turn, this suggests that women may select into matches with more rents, which also acts to reduce their cyclicality. Women may face greater volatility in their non-market activities, for instance, related to their typically greater contributions to child birth, child rearing, and other non-market activities. This may act to insulate women from the business cycle if women sort into employment matches that provide rents by providing non-pecuniary factors (e.g., flexible hours) that complement home production. But we leave to future work the fuller challenge of understanding how men and women differ.

Acknowledgments

We thank Evgenia Dechter for her excellent research assistance. We particularly thank Mark Aguiar, Iourii Manovskii, Valerie Ramey and Richard Rogerson for their helpful suggestions. Chang gratefully acknowledges support from the WCU (R33-10005) Program through the Korea Research Foundation.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.jmoneco.2012.01.001.

References

Aguiar, M., Hurst, E., 2009. Deconstructing Lifecycle Expenditure, University of Rochester, manuscript.
Anderson, P.M., Meyer, B.D., 1997. Unemployment insurance takeup rates and the after-tax value of benefits. Quarterly Journal of Economics 112, 913–937.
Costain, J.S., Reiter, M., 2008. Business cycles, unemployment insurance, and the calibration of matching models. Journal of Economic Dynamics and Control 32 (4), 1120–1155.
Fujita, S., Nekarda, C.J., Ramey, G., 2007. The Cyclicality of Worker Flows: New Evidence from the SIPP. Federal Reserve Bank of Philadelphia, Research Department Working Paper No. 07-5, February.
Guerrieri, V., Shimer, R., and Wright, R., 2009. Adverse Selection in Competitive Search Equilibrium, University of Chicago, manuscript.
Hagedorn, M., Manovskii, I., 2008. The cyclical behavior of equilibrium unemployment and vacancies revisited. American Economic Review 98 (4), 1692–1706.
Hall, R.E., 2005. Employment fluctuations with equilibrium wage stickiness. American Economic Review 95 (1), 50–65.
Hosios, A.J., 1990. On the efficiency of matching and related models of search and unemployment. Review of Economic Studies 57 (2), 279–298.
Mortensen, D., Nagypal, E., 2007. More on unemployment and vacancy fluctuations. Review of Economic Dynamics 10 (3), 327–347.
Mortensen, D., Pissarides, C., 1994. Job creation and destruction in the theory of unemployment. Review of Economic Studies 61 (3), 397–415.
Mueller, A., 2010. Separations, Sorting, and Cyclical Unemployment, Stockholm University, manuscript.

Pries, M., 2004. Persistence of employment fluctuations: a model of recurring job loss. Review of Economic Studies 71 (1), 193–215.
Sahin, A., Song, J., Hobijn, B., 2010. The unemployment gender gap during the 2007 recession. Current Issues in Economics and Finance 16 (2), 1–7 (Federal Reserve Bank of New York).
Shin, D., 2000. The consequences of rigid wages in search models. Economic Inquiry 38 (4), 641–650.
Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. American Economic Review 95 (1), 25–49.
Topel, R.H., Ward, M., 1992. Job mobility and the careers of young men. Quarterly Journal of Economics 107, 439–479.